

# The Frequency Confound in Language-Model Surprisal and Metaphor Novelty

Omar Momen<sup>1</sup> and Sina Zarriess<sup>2</sup>

CRC 1646 – Linguistic Creativity in Communication  
Bielefeld University, Germany  
{omar.hassan, sina.zarriess}@uni-bielefeld.de

## Abstract

Language-model (LM) surprisal is widely used as a proxy for contextual predictability and has been reported to correlate with metaphor novelty judgments. However, surprisal is tightly intertwined with lexical frequency. We explore this interaction on metaphor novelty ratings using two different word frequency measures. We analyse surprisal estimates from eight Pythia model sizes and 154 training checkpoints. Across settings, word frequency is a stronger predictor of metaphor novelty than surprisal. Across training stages, the surprisal–novelty association peaks at an early stage and then falls again, mirroring a similarly timed increase in the surprisal–frequency association. These results suggest that the often-reported optimal LM surprisal settings may incorrectly associate contextual predictability with metaphor novelty and processing difficulty, whereas lexical frequency may be the major underlying factor. The experiment resources are publicly available.<sup>1</sup>

## 1 Introduction

Conceptual Metaphor Theory (Lakoff and Johnson, 1980) argues that a metaphor is not only a linguistic phenomenon but also a cognitive mechanism through which abstract concepts are mapped into more concrete domains. Many of such metaphorical mappings are highly conventionalised, e.g., ARGUMENT → WAR, TIME → MONEY, or LIFE → JOURNEY.

But metaphorical mappings can vary in their degree of novelty: whereas we “live by” many conventionalised metaphors in everyday language, novel metaphors can stand out as instances of linguistic creativity. For example, the metaphorical mappings of WATER in the sentence: “The arrested

water shone and danced.”<sup>2</sup> are considered relatively novel. Novel (creative) metaphors introduce less familiar mappings and require greater interpretive effort to understand (Lai et al., 2009; Cardillo et al., 2012; Philip, 2016). Commonly, metaphor novelty is measured by human ratings (Do Dinh et al., 2018).

Recent studies have shown that metaphor novelty ratings by humans correlate significantly with both lexical frequency (Do Dinh et al., 2018; Reimann and Scheffler, 2024) and LM surprisal (Momen et al., 2026). Interestingly, Momen et al. (2026) report that, within the same model family, surprisal estimates from smaller LMs consistently correlate more strongly with metaphor novelty ratings than those from larger variants. This mirrors the inverse scaling effect observed in studies of reading times, where smaller models often provide better fits to human reading behaviour (Oh and Schuler, 2023a,b; de Varda and Marelli, 2023).

Lexical frequency has been proposed as a potential explanation for the inverse scaling effect in reading time studies. Oh et al. (2024) argue that larger LMs are better than humans at predicting low-frequency words, which weakens the correlation between their surprisal estimates and human reading-time data. Also, Opedal et al. (2024) report that non-contextual frequency has a stronger effect on predicting reading time than surprisal.

Surprisal Theory (Hale, 2001; Levy, 2008) is a common approach to studying processing difficulties in humans and LMs. The theory predicts that processing difficulty increases as a word becomes less predictable in context. Recently, the consistent inverse scaling pattern reported in studies of different phenomena such as reading time or metaphor novelty (Oh and Schuler, 2023a; de Varda and Marelli, 2023; Momen et al., 2026) has mo-

<sup>1</sup>Data and code: [https://github.com/OmarMomen14/surprisal\\_frequency\\_novelty](https://github.com/OmarMomen14/surprisal_frequency_novelty)

<sup>2</sup>Example from the VU Amsterdam Metaphor Corpus (VUAMC) (Steen et al., 2010), and the sentence is traced back to the novel “Still Life” by A. S. Byatt, 1985.

tivated an essential assumption that smaller LMs or shorter pretraining provide more human-aligned surprisal estimates. However, whether this pattern reflects genuinely more human-like contextual prediction or confounding lexical properties remains unclear.

In this paper, we analyse the associations between the three quantities: “Metaphor Novelty Ratings”, “Lexical Frequency” and “LM Surprisal”. We propose two different methods for computing lexical frequency. Then, we compute LM surprisal estimates across (i) eight Pythia model sizes and (ii) 154 pretraining checkpoints. Our results show that word frequency is a substantially stronger predictor of novelty than surprisal across settings, and that the configurations where surprisal performs “best” are also those where it aligns most closely with frequency. We therefore caution against interpreting strong early-training or small-model surprisal–novelty associations as straightforward support for surprisal-based accounts of metaphor novelty, and we argue that progress requires clearer theoretical and methodological separation between contextual predictability and lexical-frequency effects.

## 2 Data & Methods

This Section describes our experimental set-up to measure the association between word frequency and metaphor novelty on the one hand and LM surprisal on the other hand.

### 2.1 Dataset

We base our study on the VU Amsterdam Metaphor Corpus (VUAMC) (Steen et al., 2010). Every word in VUAMC is annotated as either a metaphoric word or not. Building on VUAMC, Do Dinh et al. (2018) collected crowd-sourced metaphor novelty ratings for the 15,155 metaphoric content words in VUAMC and converted them into continuous scores in the range (-1, +1), where -1 denotes the most conventional and +1 the most novel. They additionally binarised these scores using a 0.5 threshold, resulting in 353 metaphors labelled as novel out of the 15,155 content metaphoric words (see Appendix A). In our experiment, we use these 15,155 instances, each consisting of a single sentence, a target word within the sentence (a content metaphoric word), and an associated metaphor novelty score.

### 2.2 Model Suite

To examine the effects of model scale and pretraining progress (data/steps), we use the Pythia model suite (Biderman et al., 2023). Pythia consists of decoder-only causal LMs at 8 sizes (70M–12B parameters), all trained on the same 300B-token pretraining corpus<sup>3</sup> in the same order. For each model, Pythia provides 154 intermediate checkpoints saved every 1,000 training steps (corresponds to additional  $\approx 2\text{M}$  tokens seen during these steps), and denser early checkpoints at steps  $\{1, 2, 4, 8, 16, 32, 64, 128, 256, 512\}$ . The exact pretraining sequences seen at each checkpoint can be reconstructed using available scripts.

### 2.3 Surprisal

For causal LMs, surprisal, computed for a word  $w_i$  within a sequence  $W_{0:n}$  is  $\text{Surprisal}(w_i) = -\log p(w_i | w_{<i})$ <sup>4</sup>. In our experiment, we measure word-level surprisal for metaphoric target word(s) in their sentence-level context by running an independent, teacher-forced forward pass per sentence and recording the target word surprisal. To map token probabilities to a target word, we locate the word’s character offsets in the sentence and sum token-level surprisals over the minimal token span in the sequence’s tokenisation that covers these offsets. We additionally apply word-probability corrections for leading-whitespace tokenisation confounds (Pimentel and Meister, 2024), and we prepend a BOS token so surprisal is defined even when the target is the first word in a sentence.

### 2.4 Word Frequency

Estimating word frequency is non-trivial because frequency is not directly observed as a property of a word, but estimated from a reference corpus. Different corpora approximate different kinds of language exposure, and there is no single corpus that can be assumed to represent the linguistic experience of all speakers. To tackle this problem, we utilise two different estimates of word frequency reflecting both human and LM lexical memory. In both estimates, we compute the *negative log frequency* of each target metaphoric word for easier numerical comparability with surprisal.

<sup>3</sup>The Pile (Gao et al., 2020)

<sup>4</sup>We use log of base  $e$  for all log computations in our study.

1. **Negative Log Frequency in General Language Use:** We compute the negative log frequency of each target metaphoric word using the Python library *wordfreq*<sup>5</sup>, which provides corpus-independent frequency estimates aggregated from multiple large-scale sources, rather than deriving counts from a single corpus. We treat this as an estimate of word frequency for an “average English speaker”, and hereafter denote it as **NLF-Human**.
2. **Negative Log Frequency in Pythia’s Pretraining Data:** For each target metaphoric word, we tokenise its sentence using Pythia’s tokeniser and, at each checkpoint, we count occurrences of the target word’s subtoken sequence in the pretraining tokens seen up to that checkpoint. We treat the negative log of these frequencies as an LM checkpoint-specific estimate of word frequency. Hereafter, we denote this estimate as **NLF-LM**.

## 2.5 Experiment

We compute surprisal using all 8 Pythia model sizes and at each of the 154 checkpoints of Pythia-70M. We likewise compute NLF-LM at each of these checkpoints, and NLF-Human (once) per metaphoric word. We quantify surprisal–novelty and frequency–novelty associations using Spearman’s ( $\rho$ ) and Pearson’s ( $r$ ) correlation coefficients, and we report the Area Under the ROC Curve (AUC) as an estimate of discriminating binarised novel metaphors using surprisal or frequency as a predictor. Additionally, Surprisal–frequency associations are estimated using Spearman’s ( $\rho$ ) correlation.

## 3 Results

Figures 1–4 visually illustrate the results across model sizes and checkpoints. All detailed numerical results are provided in Appendix B.

### 3.1 Associations with Metaphor Novelty

**Model Scale:** Figure 1 shows surprisal–novelty and frequency–novelty association across Pythia model sizes. Here, NLF-LM is computed at the final checkpoint (300B tokens), and is therefore identical across sizes. Overall, frequency has a clearly stronger association with novelty than surprisal, with NLF-Human yielding slightly higher

estimates ( $\rho = .66, r = .66, AUC = .90$ ) than NLF-LM ( $\rho = .63, r = .60, AUC = .90$ ). We also observe a consistent negative effect of model scale on the surprisal–novelty association.


**Pretraining Progress:** Figure 2 reports associations across the 154 checkpoints of Pythia-70M. Despite the change of NLF-LM across checkpoints, NLF-LM–novelty association remains almost the same across checkpoints, except for a small deviation at the earliest steps (1–4). In contrast, the surprisal–novelty association is very weak in the first checkpoints, then rises sharply after 64 training steps (134M tokens), and peaks after 128 steps (268M tokens), where it approaches the frequency estimates ( $\rho = .62, AUC = .90$ ). After this peak, the surprisal–novelty association converges to ( $\rho = .45, AUC = .83$ ). Yet, surprisal never reaches the same strength of association with novelty as frequency does.

### 3.2 Correlations between Surprisal and Frequency


**Model Scale:** Figure 3 reports frequency–surprisal Spearman correlations across Pythia model sizes. Here, both NLF-Human and NLF-LM are fixed across sizes, while surprisal changes. Across sizes, surprisal shows moderate correlations with both NLF estimates ( $\rho \in [.40, .61]$ ). Correlations decrease with model size, mirroring the negative scale effect observed for associations with novelty (Figure 1), suggesting that larger models’ surprisal diverges more from frequency estimates. Overall, surprisal here correlates slightly more with NLF-LM (max  $\rho = .61$ ) than with NLF-Human (max  $\rho = .57$ ).

**Pretraining Progress:** Figure 4 reports frequency–surprisal correlations across the 154 checkpoints of Pythia-70M. Here, NLF-Human is fixed, while NLF-LM and surprisal vary with checkpoint. The correlation pattern of surprisal to frequency closely matches the trends of surprisal’s association to novelty in Figure 2: here surprisal has a weak correlation to frequency at the earliest checkpoints, then the correlation rise sharply after 64 training steps (134M tokens), and peak after 128 steps (268M tokens), reaching  $\rho = .95$  with NLF-LM and  $\rho = .89$  with NLF-Human. Correlations then gradually converge across subsequent checkpoints, with moderate strength  $\rho \approx .60$ .


<sup>5</sup><https://pypi.org/project/wordfreq/>




**Figure 1:** Effect of model size on associations between Metaphor Novelty Scores and Surprisal (**solid**); Negative Log Word Frequency in general language use (NLF-Human) (**dash**); and Negative Log Word Frequency in Pythia’s pretraining data (NLF-LM) (**dots**). Blue lines track Spearman correlation, and red lines track AUC to detect novel metaphors ( $score \geq 0.5$ ).



**Figure 2:** Effect of pretraining data/steps for Pythia-70M on associations between Metaphor Novelty Scores and Surprisal (**solid**); Negative Log Word Frequency in general language use (NLF-Human) (**dash**); and Negative Log Word Frequency in Pythia’s pretraining data (NLF-LM) (**dots**). Blue lines track Spearman correlation, and red lines track AUC.



**Figure 3:** Effect of model scale on correlation between Surprisal and Frequency. NLF-Human (**dash**); and NLF-LM (**dots**).



**Figure 4:** Effect of pretraining data/steps of Pythia-70M on correlation between Surprisal and Frequency. NLF-Human (**dash**); and NLF-LM (**dots**).

## 4 Discussion

**Word Frequency:** Our results agree with previous work (Do Dinh et al., 2018; Reimann and Scheffler, 2024) showing that lexical frequency is strongly associated with metaphor novelty scores. Additionally, we show that frequency–novelty association is substantially stronger than surprisal–novelty associations across all LM settings in our experiment. Notably, frequency correlates more strongly with novelty than it does with surprisal itself. We further

observe that NLF-Human aligns slightly more with novelty (human-based) than NLF-LM, whereas NLF-LM aligns slightly more with surprisal (LM-based) than NLF-Human. Overall, however, differences between the two frequency estimates are small and do not alter overall trends, suggesting that estimating frequencies from relatively small amounts of data is sufficient (at least for our task) when large-scale estimation is expensive.

**Inverse Scale Effect:** The surprisal–novelty association decreases with model size, replicating prior results on the same dataset across other model families (Momen et al., 2026), and on datasets of reading time (Oh and Schuler, 2023b). We additionally show that the same negative scaling trend holds for correlations between surprisal and frequency: as model size increases, surprisal becomes less aligned with frequency.

**Pretraining Amount:** The strongest association between surprisal and novelty, and between surprisal and frequency, occurs early—after 128 pretraining updates for Pythia-70M ( $\approx 268$ M tokens seen)—and additional training weakens this association. A qualitatively similar non-monotonic effect has been reported for reading times: surprisal predicts reading time best at an intermediate pretraining amount (about 2B tokens), after which further pretraining reduces predictive power (Oh and Schuler, 2023a). The close similarity between the checkpoint trends for surprisal–novelty (Figure 2) and surprisal–frequency (Figure 4) highlights the extent to which word frequency can confound surprisal-based analyses of linguistic and psycholinguistic targets.

**Surprisal:** Surprisal achieves its strongest association with metaphor novelty when computed from the smallest model (70M) and at a relatively early training stage (128 steps;  $\approx 268$ M tokens). However, these are also the settings in which surprisal is most closely aligned with word frequency. We therefore caution against treating these “best” association values as direct evidence for surprisal-as-predictability accounts of metaphor novelty (and processing difficulty in general): in these settings, surprisal may primarily reflect lexical frequency, and possibly more than contextual predictability. We think that larger models or extensive pretraining do not produce intrinsically poor estimates of predictability. And, we call for further efforts to develop more accurate human ratings that reflect clear theories of the novelty/originality dimension in creativity (and potentially processing difficulty more broadly) and to clarify how these constructs relate to surprisal-based accounts.

## 5 Conclusion

We analysed the associations between metaphor novelty ratings, LM surprisal and lexical frequency across different model sizes and pretraining stages.

The results mainly show that lexical frequency is a stronger predictor of metaphor novelty than LM surprisal. We further find that the strongest correlations between LM surprisal and metaphor novelty occur only under settings in which surprisal is also strongly correlated with lexical frequency. We conclude that relevant studies should be careful not to interpret strong surprisal–novelty associations at small or early-training LM settings as straightforward evidence for surprisal-as-predictability accounts of metaphor novelty (or processing effort).

## Limitations

Due to computational constraints, we do not compute surprisal for intermediate checkpoints of the larger Pythia models, and we restrict the checkpoint-level analysis to the smallest variant (Pythia-70M). Although this choice is consistent with our model-scale findings (smaller models yield stronger associations), evaluating intermediate checkpoints for larger models remains necessary to verify whether the observed training-dynamics trends hold across model sizes.

## Acknowledgments

This research has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – CRC-1646, project number 512393437, project A05.

## References

- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. 2023. *Pythia: A suite for analyzing large language models across training and scaling*. Preprint, arXiv:2304.01373.
- Eileen R. Cardillo, Christine E. Watson, Gwenda L. Schmidt, Alexander Kranjec, and Anjan Chatterjee. 2012. *From novel to familiar: Tuning the brain for metaphors*. *NeuroImage*, 59(4):3212–3221.
- Andrea de Varda and Marco Marelli. 2023. *Scaling in cognitive modelling: a multilingual approach to human reading times*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 139–149, Toronto, Canada. Association for Computational Linguistics.
- Erik-Lân Do Dinh, Hannah Wieland, and Iryna Gurevych. 2018. *Weeding out conventionalized metaphors: A corpus of novel metaphor annotations*.

- In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1424, Brussels, Belgium. Association for Computational Linguistics.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. [The pile: An 800gb dataset of diverse text for language modeling](#). Preprint, arXiv:2101.00027.
- John Hale. 2001. [A probabilistic earley parser as a psycholinguistic model](#). In *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics on Language Technologies*, NAACL '01, page 1–8, USA. Association for Computational Linguistics.
- Vicky Tzuyin Lai, Tim Curran, and Lise Menn. 2009. [Comprehending conventional and novel metaphors: An erp study](#). *Brain Research*, 1284:145–155.
- George Lakoff and Mark Johnson. 1980. *Metaphors We Live By*. University of Chicago Press, Chicago, IL.
- Roger Levy. 2008. Expectation-based syntactic comprehension. *Cognition*, 106(3):1126–1177.
- Omar Momen, Emilie Sitter, Berenike Herrmann, and Sina Zarrieß. 2026. [Surprisal and metaphor novelty judgments: Moderate correlations and divergent scaling effects revealed by corpus-based and synthetic datasets](#). In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8113–8127, Rabat, Morocco. Association for Computational Linguistics.
- Byung-Doh Oh and William Schuler. 2023a. [Transformer-based language model surprisal predicts human reading times best with about two billion training tokens](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1915–1921, Singapore. Association for Computational Linguistics.
- Byung-Doh Oh and William Schuler. 2023b. [Why does surprisal from larger transformer-based language models provide a poorer fit to human reading times?](#) *Transactions of the Association for Computational Linguistics*, 11:336–350.
- Byung-Doh Oh, Shisen Yue, and William Schuler. 2024. [Frequency explains the inverse correlation of large language models' size, training data amount, and surprisal's fit to reading times](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2644–2663, St. Julian's, Malta. Association for Computational Linguistics.
- Andreas Opedal, Eleanor Chodroff, Ryan Cotterell, and Ethan Wilcox. 2024. [On the role of context in reading time prediction](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3042–3058, Miami, Florida, USA. Association for Computational Linguistics.
- Gill Philip. 2016. [Conventional and novel metaphors in language](#). In Elena Semino and Zsófia Demjén, editors, *The Routledge Handbook of Metaphor and Language*. Routledge, London / New York. “Chapter 15: Conventional and Novel Metaphors in Language”.
- Tiago Pimentel and Clara Meister. 2024. [How to compute the probability of a word](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18358–18375, Miami, Florida, USA. Association for Computational Linguistics.
- Sebastian Reimann and Tatjana Scheffler. 2024. [When is a metaphor actually novel? annotating metaphor novelty in the context of automatic metaphor detection](#). In *Proceedings of the 18th Linguistic Annotation Workshop (LAW-XVIII)*, pages 87–97, St. Julians, Malta. Association for Computational Linguistics.
- G.J. Steen, A.G. Dorst, J.B. Herrmann, A.A. Kaal, T. Krennmayr, and T. Pasma. 2010. *A method for linguistic metaphor identification. From MIP to MIPVU*. Number 14 in *Converging Evidence in Language and Communication Research*. John Benjamins.

## A Dataset Statistics

In Table 1, we demonstrate the statistics of the dataset used in our experiment. In Table 2, we list 16 examples from the VUAMC with their associated novelty ratings from Do Dinh et al. (2018).

## B Numerical Results

Detailed numerical results of our study are listed in Tables 3, 4, 5 and 6.

Genre	# Metaphors	$L_{sent}$	Novelty Score	# Novel
		mean   std.	mean   std.	$\geq 0.5$
Fiction	3170	26.0   16.5	-.005   .271	94
News	4712	29.9   14.2	.000   .257	132
Academic	5499	34.9   16.0	.003   .239	102
Conversation	1774	17.5   15.9	-.000   .236	25
All	15155	29.4   16.5	.000   .251	353

**Table 1:** Distributions and statistics of the dataset under study. # Metaphors is the number of metaphor words,  $L_{sent}$  values are the statistics of sentences lengths in words. Novelty Score values are the normalised continuous novelty ratings, # Novel is the number of novel metaphors (Novelty Score  $\geq 0.5$ ).

Genre	Sentence	Novelty Score	Label
Fiction	1. ‘ Tell him I am very sorry, but I must <b>fill</b> the quota. ’	-0.441	conventional
	2. Adam might have escaped the file memories for years, <b>suppressed</b> them and jerked violently <14> by those events.	0.531	novel
	3. It was an excitement that <11> and I had long dreamed of that scatter of tiny, <b>magically</b> named islands strewn across one third of a globe.	0.278	conventional
	4. The seemingly random and <11> designed to disguise a boat’s shape from the <b>prying</b> eyes of U-Boat captains, so it <10> in the Bahamas.	0.588	novel
News	5. One Mr Clarke can not duck away from if he wants to <b>avoid</b> a second Winter of Discontent	-0.094	conventional
	6. This was conveniently <b>encapsulated</b> in the first try.	0.500	novel
	7. Thrusts of resistance ( mass demonstrations, resignations, tax rebellions, etc ) would come in <b>crests</b> .	0.382	conventional
	8. Travel: A <b>pilgrimage</b> sans progress Elisabeth de Stroumillo potters round Poitou	0.514	novel
Academic	9. The Tehuana dress is by no means the most decorative variant or the <b>closest</b> to pre-Hispanic forms of clothing.	-0.194	conventional
	10. Interwoven with these images are subtler references to the <b>metaphorical</b> borderlines which separate Latin American <5> and North America.	0.529	novel
	11. This is often linked with a supposed denunciatory effect — the idea that the mandatory life sentence <b>denounces</b> murder as emphatically as possible <18> this crime.	0.294	conventional
	12. He certainly held deep convictions as to the <9>, but at least a part of his apparent <b>hostility</b> was assumed for the occasion, a hard <7> in the end.	0.514	novel
Conversation	13. Me dad said he’s had enough Well, we were debating whether to <b>give</b> it to you or not.	-0.633	conventional
	14. <b>Struggled</b> with it a little	0.552	novel
	15. That’s an old <b>trick</b> .	0.310	conventional
	16. Can you <b>sort</b> erm, madame out?	0.567	novel

**Table 2:** Examples from the VU Amsterdam Metaphor Corpus (VUAMC). The metaphor word is in **boldface** within sentences. For simpler presentations, we remove some words from long sentences and replace them with a tag of the number of words removed, e.g. <11>. **Novelty Score** is the Do Dinh et al. (2018) normalised human rating, and **Label** is the binary novelty label based on the 0.5 threshold. Contrastive examples are picked randomly from the dataset for each genre to illustrate the differences between conventional and novel instances according to the “Novelty Score”.

Model	Pearson ( $r$ )	Spearman ( $\rho$ )	AUC
NLF-Human	<b>.656</b>	<b>.664</b>	<b>.904</b>
NLF-LM	<b>.599</b>	<b>.633</b>	<b>.899</b>
Pythia-70M	<b>.448</b>	<b>.453</b>	.832
Pythia-160M	.426	.425	<b>.833</b>
Pythia-410M	.382	.374	.792
Pythia-1B	.371	.359	.782
Pythia-1.4B	.357	.344	.771
Pythia-2.8B	.351	.336	.766
Pythia-6.9B	.338	.322	.758
Pythia-12B	.330	.312	.752

**Table 3:** Pearson’s  $r$  and Spearman’s  $\rho$  Correlation and AUC estimates between **Metaphor Novelty Scores** and **Surprisal; Negative Log Word Frequency in general language use (NLF-Human)**; and **Negative Log Word Frequency in Pythia’s pretraining data (NLF-LM)** across different model sizes. All reported estimates are significant at the 0.001 level.

Model	NLF-Human ( $\rho$ )	NLF-LM ( $\rho$ )
Pythia-70M	<b>.572</b>	<b>.606</b>
Pythia-160M	.533	.573
Pythia-410M	.467	.507
Pythia-1B	.449	.489
Pythia-1.4B	.434	.474
Pythia-2.8B	.421	.462
Pythia-6.9B	.404	.445
Pythia-12B	.396	.438

**Table 4:** Spearman Correlation estimates between **Surprisal** and **Negative Log Word Frequency in general language use (NLF-Human)**; and **Negative Log Word Frequency in Pythia’s pretraining data (NLF-LM)** across different model sizes. All reported estimates are significant at the 0.001 level.

# Steps	# Pretraining Tokens	Surprisal		NLF-LM	
		$\rho$	AUC	$\rho$	AUC
1	2M	.221	.799	.626	.891
2	4M	.221	.799	.631	.894
4	8M	.221	.799	.631	.896
8	17M	.224	.800	.631	.897
16	34M	.226	.799	.631	.898
32	67M	.250	.804	.632	.898
64	134M	.435	.858	.632	<b>.899</b>
128	268M	<b>.618</b>	<b>.895</b>	.632	.899
256	537M	.574	.893	<b>.633</b>	.899
512	1B	.554	.888	.633	.899
1,000	2B	.524	.873	.633	.899
2,000	4B	.473	.849	.633	.899
3,000	6B	.475	.846	.633	.899
4,000	8B	.458	.842	.633	.899
5,000	10B	.458	.842	.633	.899
6,000	13B	.459	.841	.633	.899
7,000	15B	.457	.836	.633	.899
8,000	17B	.453	.833	.633	.899
9,000	19B	.451	.833	.633	.899
12,000	25B	.459	.838	.633	.899
17,000	36B	.448	.832	.633	.899
22,000	46B	.447	.831	.633	.899
27,000	57B	.446	.836	.633	.899
32,000	67B	.453	.835	.633	.899
37,000	78B	.446	.839	.633	.899
42,000	88B	.450	.843	.633	.899
47,000	99B	.445	.833	.633	.899
52,000	109B	.444	.829	.633	.899
57,000	120B	.443	.827	.633	.899
62,000	130B	.442	.831	.633	.899
67,000	141B	.446	.831	.633	.899
72,000	151B	.441	.830	.633	.899
77,000	161B	.447	.828	.633	.899
82,000	172B	.445	.824	.633	.899
87,000	182B	.440	.826	.633	.899
92,000	193B	.442	.831	.633	.899
97,000	203B	.443	.826	.633	.899
102,000	214B	.449	.828	.633	.899
107,000	224B	.448	.829	.633	.899
112,000	235B	.453	.833	.633	.899
117,000	245B	.447	.828	.633	.899
122,000	256B	.455	.835	.633	.899
127,000	266B	.451	.830	.633	.899
132,000	277B	.451	.835	.633	.899
137,000	287B	.448	.829	.633	.899
143,000	300B	.453	.832	.633	.899

**Table 5:** Spearman Correlation and AUC estimates between **Metaphor Novelty Scores** and **Surprisal**; and **Negative Log Word Frequency in Pythia’s pretraining data (NLF-LM)** across pretraining steps of Pythia-70M, reporting the amount of pretraining data tokens seen at each step.

# Steps	# Pretraining Tokens	NLF-Human ( $\rho$ )	NLF-LM ( $\rho$ )
1	2M	.286	.302
2	4M	.286	.306
4	8M	.286	.309
8	17M	.292	.316
16	34M	.305	.332
32	67M	.352	.388
64	134M	.623	.679
128	268M	<b>.885</b>	<b>.953</b>
256	537M	.792	.866
512	1B	.730	.778
1,000	2B	.679	.705
2,000	4B	.597	.629
3,000	6B	.601	.634
4,000	8B	.580	.615
5,000	10B	.587	.623
6,000	13B	.583	.609
7,000	15B	.579	.614
8,000	17B	.578	.606
9,000	19B	.572	.609
12,000	25B	.583	.614
17,000	36B	.572	.605
22,000	46B	.566	.595
27,000	57B	.566	.599
32,000	67B	.578	.612
37,000	78B	.563	.602
42,000	88B	.570	.604
47,000	99B	.564	.600
52,000	109B	.562	.598
57,000	120B	.559	.595
62,000	130B	.557	.591
67,000	141B	.560	.595
72,000	151B	.559	.597
77,000	161B	.565	.601
82,000	172B	.563	.597
87,000	182B	.557	.595
92,000	193B	.557	.592
97,000	203B	.557	.594
102,000	214B	.565	.597
107,000	224B	.563	.597
112,000	235B	.571	.604
117,000	245B	.562	.595
122,000	256B	.572	.607
127,000	266B	.568	.602
132,000	277B	.567	.602
137,000	287B	.563	.600
143,000	300B	.572	.606

**Table 6:** Spearman Correlation estimates between **Surprisal** and **Negative Log Word Frequency in general corpora**; and **Negative Log Word Frequency in pretraining corpora** across pretraining steps of Pythia-70M, reporting the amount of pretraining data tokens seen at each step.